unpatentable over Searby. Applicant respectfully traverses these grounds of rejection and requests reconsideration for the following reasons.

Claim 1

Claim 1 is directed to a method of reading data from a RAID array of disk drives. In particular, claim 1 focuses on one aspect of the present invention, namely reconstruction of erroneous data during a read operation. Data errors can arise, for example, from corruption or failure of one of the disk drives in the array. One of the steps in the claimed method calls for shifting the read data through a shift register -- an element the Examiner contends corresponds to the "delay shift register" 57 shown in Searby at figure 3. However, "a claim is anticipated only if each and every element as set forth in the claim is found, either expressly or inherently described, in a single prior art reference," and the elements must be arranged as required by the claim. MPEP 2131 (Rev.1, Feb. 2000).

The circuitry taught by Searby is plainly distinguished in that the delay shift register 57 of figure 3 is disposed between the RAM buffer memory 40 and the external data bus (register 50) which is connected to the host. This arrangement is markedly different from the architecture of the present invention, as illustrated in applicant's figure 9, which shows that applicant's shift register 902 is disposed between the disk drive data ports (900) and the buffer memory (in figure 9, "TO CACHE"). As described in the text:

> "The processor will use the circuitry of Figure 9 to reconstruct the data stored in the bad drive during the read operation as follows. In Figure 9, word-serial read data from the disk drives is received via bus 900 into a data pipeline 902.... Data flows out of the pipeline 902 via path 952 into a first input to multiplexer 920 and then to the cache or buffer memory."

Specification, page 25, lines 20-28. In Searby, there is no shift register in between the data port 36 and the RAM buffer 40.

Claim 1 is amended to clarify these distinctions. For example, the amended claim calls for "providing a series of registers forming a common pipeline disposed in between the disk drive data ports and the buffer memory data port." The claim is also amended to clarify that the corrected read data is stored in the buffer memory, whereas in the

architecture taught by Searby, potentially erroneous or corrupted data is stored in the buffer memory. The correction takes place both physically and temporally in between the buffer memory and the external host.

The claimed architecture of the present invention provides several advantages in this regard. The present invention is intended primarily to address the general RAID market place rather than the video editing application described by Searby. The video editing application as described by Searby included the use of 20 drives and, as Searby observes, the corresponding cost for the redundant drive and related hardware would add only five percent to the cost of the overall system. The present invention is directed to a high volume RAID application where the entry-level system might start out as a 2-drive mirroring system and then be expanded to a 3-drive RAID 3 system. For these types of applications, introducing the redundancy logic *between the drive interface and the buffer memory*, as described in claim 1, provides the following advantages.

1. Memory Bandwidth: While RAM bandwidths are much greater than disk drive bandwidths, the total requirements for an array of drives and concurrent host transfers make the buffer bandwidth a critical resource. For a 3-drive application consisting of one redundant drive and two data drives, processing redundant data between the buffer and the host system (as taught by Searby) adds 50% to the buffer bandwidth requirements as 50% more data must be transferred in and out of the buffer memory for a given data throughput. In contradistinction, according to the present invention, no redundant data is stored in the buffer memory, since the redundancy and data correction logic is deployed between the buffer memory and the drives.

2. Memory Capacity: If the buffer memory is going to provide some level of data caching, 50% more memory would be required to cache data plus redundancy for a 3-drive RAID system.

3. Memory Scalability: The design and management of the buffer memory grows increasing complex as scalability is considered. For a 3-drive system, 1/3 of the memory is used to store redundant data and must be accessed concurrently with the data during host transfers. For a 5-drive system, 1/5 of the memory stores redundant data and must be accessed concurrently with host transfers. These inefficiencies are avoided by

the methodology described in applicant's amended claims 1. For these reasons, claim 1 should be allowed.

Claim 2

Claim 2 is directed to a disk array controller and calls for "redundant data operating means disposed along the drive data bus for forming redundant drive data on the fly as data passes from the buffer memory to the drives during a disk write operation." (Emphasis added.) As explained above, this arrangement is in contradistinction to the prior art which discloses redundant data operating means disposed between the buffer memory and host. The Examiner points out disk interfaces 25-28, which could be considered "disk drive interface means" but Searby's redundancy logic (55,56,57 in figure 3) is not between the buffer memory and the drives and therefore it cannot be used for "forming redundant drive data *on the fly* as data passes from the buffer memory to the drives during a disk write operation," as required by claim 2. Rather, in Searby, redundant data is formed and stored in the buffer 40 before it is written to the drives. Moreover, the term "drive data bus" refers to a data bus interconnecting the drives and the buffer memory. For example, see the multiplexed drive data bus 510 in figure 5. Thus the redundancy logic shown in Searby does not anticipate "redundant data operating means disposed along the drive data bus" as per claim 2. For these reasons claim 2 should be allowed.

Claims 3-4 and 7-9

Claims 3 and 4 depend from claim 2 and add further detail to the novel architecture described above. For example, in claim 3, the redundant data operating means is described as including "a multiplexer having a first input coupled to the buffer memory port to receive write data" and "the multiplexer output coupled to the drive data bus for interfacing to the array of disk drives...". Thus it is clear in accordance with the language of claim 3 that the redundant data operating means is disposed in between the buffer memory and the disk drive data ports.

Applicant also respectfully disagrees with the Examiner's contention that the parity generator 55 of Searby's figure 3 is "disposed along the drive data bus." The "drive data bus" as noted above serves to interconnect the disk drive data ports and the buffer memory, whereas the parity generator 55 is disposed along a data path *between the*

buffer memory and the host; it is segregated from the disk drive. With respect to claim 4, applicant would also point out that the data generator 56 and delay shift register 57 likewise are not disposed along the disk drive data bus as that term is used on the claim. Locating the redundancy and error correction logic between the drives and the buffer memory provides important advantages as described above with reference to claim 1. For these reasons, claims 3 and 4 should be allowed.

Claims 6-8 and 10-19

The Examiner also rejected claims 6-8 and 10-19 under Section 102(e) as being anticipated by Searby. Applicant respectfully traverses these rejections and requests reconsideration for the following reasons. With regard to claim 6, it calls for "means disposed between the buffer memory and the data bus for generating redundant check data on the fly during execution of a disk write operation." What data bus? Claim 6 explicitly recites "a data bus interconnecting the buffer memory and the disk drive interface means," and the claim further recites that the disk drive interface means provides "for connection to a plurality of disk drives." Thus it is clear from parsing the claim language that it requires the means for generating redundant check data to be disposed between the disk drives and the buffer memory. This architecture is squarely contrary to that shown in the reference as discussed above with regard to other claims. For these reasons, claim 6 also should be allowed.

Claim 7 is similar to claim 6, except that it calls for "means disposed between the buffer memory and the drive data bus for reconstructing missing data during a read operation so that only correct read data is stored in the buffer memory." This element is contrary to the Searby reference for the reasons explained above. The architecture is different, and the resulting operation is quite different in that applicant's architecture provides for storing only correct read data in the buffer memory. Searby must correct erroneous data stored in the buffer before shipping it to the host. Claims 8 and 9 depend from claim 7 and add further detail. For these reasons, claims 7-9 should be allowed.

The Examiner characterized claims 10-19 as the corresponding method steps of the claimed apparatus. Applicant therefore respectfully traverses this ground of rejection, and requests reconsideration, based on the foregoing remarks.

Section 103(a) Rejection of claims 5 and 9

The Examiner rejected claims 5 and 9 as unpatentable over Searby. Claim 5 depends from claim 4 and thus is patentable for the reasons explained above with regard to claim 4. Claim 9 depends from claim 8 which in turn depends from claim 7, and thus claim 9 is patentable for the reasons explained above with regard to claim 7.

In view of the present amendments to the claims and the foregoing remarks, all of the pending claims should now be allowed. The Examiner is invited to telephone the undersigned if any issues remain. Email may be directed to mdstolowitz@stoel.com.

Respectfully submitted,

Michael C. Stolowitz

By: 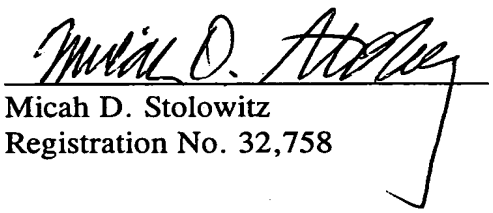

Micah D. Stolowitz
Registration No. 32,758

STOEL RIVES LLP
900 SW Fifth Avenue, Suite 2600
Portland, Oregon 97204-1268
Telephone: (503) 224-3380
Facsimile: (503) 220-2480
Attorney Docket No. 44541/2:2